Peter Boncz (CWI)

Self-organizing access and storage of correlated data: crossing the frontier of structured data management

Database systems have reached a high maturity level only if one considers them within a restricted set of usage assumptions. One of such assumptions is that the selectivity of different predicates is statistically independent, another is that all performance-critical access patterns can be predicted and elicited, and the system has time to prepare and create full indices on these. In this talk, I sketch the implications such assumptions have for modern data management problems, arguing that RDF data stores are especially hard hit by these. Consequently, I will making a case for research into adaptive storage structures, indexing methods, as well as query optimization techniques that recognize and exploit correlations. This line of thought will illustrated using recent results from the CWI database group in the areas of on-the-fly data reorganization ("cracking"), partial result caching ("recycling") as well as data correlation-aware run-time query optimization ("ROX").

Frank van Harmelen (VUA)

Knowledge Representation on the Web: what to do when success is becoming a problem

In recent years, the Semantic Web has seen rapid growth in size (many billions of facts and rules are now available) and increasing adoption in many sectors (government, publishing industry, media).
This success has brought with it a whole new set of problems: storage, querying and reasoning with billions of facts and rules that are distributed across different locations. We will argue that we must break with some well established articles of faith in classical knowledge representation in order to deal with the new challenges that come with such scale. We will also show how some of the newly developed methods have been implemented in the Large Knowledge Collider (LarKC), an infrastructure that exploits parallelization, distribution and approximation to enable Web reasoning at arbitrary scale.

Hajo Reijers (TUE)

On the Usability of Business Process Models: A Formal and Empirical Approach

Business process models are widely used artifacts to streamline the communication between professionals and to serve as blueprints for IT projects. While the formal foundations for business process models have been studied extensively, research into their usability picked up only recently.
We will present an approach to improve the readability of business process models by a sensible use of color for its elements. Both the empirical motivation for this approach and the formal basis for selecting the elements to be colored will be addressed. An implementation is provided through the publicly available WoPeD modeling tool.

Michael Wooldridge (Liverpool)

"Professor Kripke, let me introduce Professor Nash: Logic for Economic Mechanism Design"

Recent years have seen an enormous growth of interest in work at the intersection of logic and game theory. For example, researchers have used logic to make explicit the assumptions such as common knowledge that underpin many solution concepts, and have shown how solution concepts can be given a precise logical characterization. From the perspective of computer science, these links open up the possibility of applying automated verification techniques, such as model checking, to the analysis and verification of computational mechanisms. We describe our work in this area, focusing on the use of game theoretic concepts in the specification and verification of coordination mechanisms for multi-agent systems.

Arno Siebes (UU)

The Minimum Description Length Principle for Data Mining

Abstract will be available shortly